



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 12/12, 12/08	A1	(11) International Publication Number: WO 98/41923 (43) International Publication Date: 24 September 1998 (24.09.98)
(21) International Application Number: PCT/SE98/00454 (22) International Filing Date: 13 March 1998 (13.03.98) (30) Priority Data: 08/818,340 14 March 1997 (14.03.97) US (71) Applicant: TELEFONAKTIEBOLAGET LM ERICSSON (publ) [SE/SE]; S-126 25 Stockholm (SE). (72) Inventor: LIEDBERG, Per; Johan Skyttesväg 234, S-125 34 Älvsjö (SE). (74) Agent: ERICSSON TELECOM AB; IPR Management & Patent Dept., S-126 25 Stockholm (SE).		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i>

(54) Title: PENALTY-BASED CACHE STORAGE AND REPLACEMENT TECHNIQUES**(57) Abstract**

Cache data replacement techniques enable improved performance in a computer system having a central processing unit (CPU), a cache memory and a main memory, wherein the cache memory has a plurality of data items stored therein. The cache data replacement techniques include associating a priority value with each of the stored data items, wherein for each data item, the priority value is an estimate of how much CPU stall time will occur if an attempt is made to retrieve the data item from the cache memory when the data item is not stored in the cache memory. When a cache entry must be replaced, the priority values are analyzed to determine a lowest priority value. One of the data items that has the lowest priority value is selected and replaced by a replacement data item. The priority value of a data item may be determined, as a function of how many other instructions have been fetched and stored in a buffer memory between a time interval defined by initiation and completion of retrieval of the data item from the main memory, wherein execution of the other instruction is dependent on completing retrieval of the data item. In other aspects of the invention, the priority values of cache entries may periodically be lowered in order to improve the cache hit ratio, and may also be reinitialized whenever the associated data time is accessed, in order to ensure retention of valuable data items in the data cache.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

-1-

PENALTY-BASED CACHE STORAGE AND REPLACEMENT TECHNIQUES

5

BACKGROUND

The present invention relates to cache memories, and more particularly to strategies for selecting data to be replaced in a cache memory.

Processing inside a computer system may be performed by a hardware element called a central processing unit (CPU). Instructions and data for the CPU may be stored in a large, main memory. The operating speed of a CPU (i.e., the time it takes to perform one instruction) is typically very much faster than the access speed of the main memory. Consequently, the CPU may be forced to idly wait for a requested instruction or data item while the main memory cycles through one memory access operation. This idle wait time seriously degrades the effective processing speed of the CPU.

In order to address this problem, a cache memory unit is often designed into the computer system. Cache memories are well-known in the computer arts as being auxiliary memories that provide a buffering capability between a CPU and a main memory. The cache memory is typically designed to run much faster than the main memory, and to be loaded out of the main memory.

Memory devices that run at the speed of the CPU are much more expensive and physically larger than the slower devices that make up the main memory. As a result, the size of a cache memory (as measured by the number of separately addressable storage cells contained within the memory) is much smaller than the size of a main memory. Because the cache memory cannot contain all of the instructions and data stored in the main memory, the CPU occasionally requests a particular instruction or data item that is not presently stored in the cache memory. Such an occurrence is called a "cache miss", and requires that the

-2-

requested instruction or data item be retrieved from main memory, stored into the cache memory, and then supplied to the CPU. It can be seen, then, that each cache miss has the potential for making the CPU wait for as long (if not longer) than it would if the cache memory were not present.

A technique for reducing the processing speed penalty whenever a data read cache miss occurs is to make instruction execution out-of-order. This means that instructions subsequent to the one that caused the cache miss will continue to execute while the CPU is waiting for the missing data. For this strategy to work, it is necessary that execution of these subsequent instructions not be dependent on the missing data. Execution of instructions that do depend on the missing data must be held in abeyance (e.g., in queues) until the missing data becomes available. When the data does become available, all of the instructions that were dependent on this data are then executed. Out-of-order instruction execution techniques are described in William Johnson, Superscaler Microprocessor Design, 1991 (ISBN 0-13-875634-1) which is incorporated herein by reference.

Even if the out-of-order execution strategy is adopted, there will likely be branch instructions in the program whose target location is in some way conditional on the missing data. One strategy for avoiding a long delay in the instruction fetching operation of the CPU is to guess which branch will be taken, and to tentatively continue fetching and executing instructions from the guessed branch. If, when the missing data becomes available, it is found that the guess was correct, then the results of the tentative execution can be made permanent (e.g., by storing results into target memory locations). However, if an incorrect guess was made, then all of the results from instructions executed after the conditional branch instruction must be flushed, and program execution restarted

-3-

from the correct branch path. A wrong guess, therefore, causes a very high performance penalty.

This strategy can be improved by further including a branch prediction memory that stores statistics on the results of previous conditional branches in order to increase the probability of making a correct guess regarding which is the correct path for a pending conditional branch. Notwithstanding the use of this strategy, there will inevitably be branches that are mispredicted, thereby causing a high performance penalty.

Another factor that influences the effective execution speed of the CPU is the fact that when a cache miss occurs, data (or one or more instructions for the case where data and instructions share the same cache memory) must be removed from the cache memory in order to make room for the missing data item. The strategy for selecting data to be removed from the cache (called a "cache replacement strategy") can also influence the effective execution speed of the CPU because the "cast out" data may be needed at a later time, thereby causing another cache miss.

Existing cache replacement strategies have been based on maximizing the probability that a requested instruction or data item will be successfully located in the cache (called a "cache hit"). One such strategy selects for removal that data item that has been least recently used (LRU) by the executing program. The basis for this approach is the concept of temporal locality: the notion that the probability that the next address to be accessed will be the same as a recently accessed address is higher the sooner the second access occurs with respect to the first.

Other cache replacement strategies are random replacement and first-in-first-out (FIFO).

All of the above cache replacement strategies have as a goal a high cache hit ratio, usually defined as the number of times an attempted cache read is successful at obtaining the data from the cache divided by the total

-4-

number of attempted cache accesses. (A related measure is the cache miss ratio, usually defined as 1 - cache hit ratio.) However, these cache replacement strategies are deficient because they fail to take into account the effects of cache misses, which will inevitably occur.

SUMMARY

It is therefore an object of the present invention to increase the effective execution speed of a CPU by means of improved methods and apparatus for determining which entries in a cache memory should be removed to make room for a replacement cache data item.

The foregoing and other objects are achieved by utilizing a cache data replacement technique in a computer system having a central processing unit (CPU), a cache memory and a main memory, wherein the cache memory has a plurality of data items stored therein. In accordance with one aspect of the invention, the cache data replacement technique includes associating a priority value with each of the stored data items, wherein for each data item, the priority value is an estimate of how much CPU stall time will occur if an attempt is made to retrieve the data item from the cache memory when the data item is not stored in the cache memory.

In another aspect of the invention, the priority value for each of the data items is ascertained by determining how many other instructions have been fetched and stored in a buffer memory between a time interval defined by initiation and completion of retrieval of the data item from the main memory, wherein execution of the other instructions is dependent on completing retrieval of the data item.

In still another aspect of the invention, the step of determining, for each of the data items, how many other instructions have been fetched and stored in the buffer memory between the time interval further comprises adjusting

-5-

the priority value by an amount based on whether any of the other instructions is a conditional branch instruction.

In yet another aspect of the invention, the step of determining, for each of the data items, how many other
5 instructions have been fetched and stored in the buffer memory between the time interval, further comprises adjusting the priority value by an amount based on whether any of the other instructions is a non-branch instruction.

In another aspect of the invention, when a
10 replacement data item is to be stored into the cache, the cache entry to be replaced is determined by analyzing the priority values to determine a lowest priority value. Then, one of the data items that has the lowest priority value is selected and replaced by the replacement data item.

15 In yet another aspect of the invention, the hit ratio of the data cache is improved by not requiring that the priority values of associated data items remain static. Instead, each of the priority values is periodically adjusted by a predetermined amount that results in a lower
20 priority value. In this way, even those cache entries having initially high priority values may eventually become candidates for replacement.

In still another aspect of the invention, valuable data entries may be permitted to remain in the data cache
25 indefinitely by adjusting the priority value by a predetermined amount in response to a read access of the associated data item, wherein the adjustment results in a higher priority value. Alternatively, the priority value can be set to an initial value in response to a read access
30 of the associated data item. When combined with the technique of periodically reducing the priority value of data entries, this technique allows valuable cache entries to remain in the data cache so long as they continue to be accessed by the running program, and further allows them to
35 be replaced when they are no longer utilized.

-6-

BRIEF DESCRIPTION OF THE DRAWINGS

The above-described and other objects and advantages of the invention will be understood by reading the following detailed description in conjunction with the drawings in which:

FIG. 1 is a block diagram of an exemplary computer system embodying the invention;

FIG. 2 is a flow chart that depicts a technique for retrieving a data item from main memory and storing it into a data cache in accordance with one aspect of the invention;

FIG. 3 is a diagram of an exemplary embodiment of a data cache for use with the inventive cache storage and replacement techniques;

FIG. 4 is a flow chart that depicts a strategy for replacing data in a cache memory in accordance with another aspect of the invention;

FIG. 5 is a diagram of an alternative exemplary embodiment of a data cache for use with the inventive cache storage and replacement techniques; and

FIG. 6 is a flow chart that depicts a strategy for maintaining priority values in accordance with other aspects of the invention.

DETAILED DESCRIPTION

The various features of the invention will now be described with respect to the figures, in which like parts are identified with the same reference characters.

An exemplary computer system embodying the invention is shown in FIG. 1. A CPU 101 executes instructions which are stored in a main memory 107. The CPU 101 retrieves data for the instructions from an address space contained in the main memory 107. Results from instruction execution are also stored back into this same address space.

-7-

The cycle time of the main memory 107 is much slower than the execution speed of the CPU 101. In order to more quickly feed the CPU 101 with instructions, an instruction cache 103 is interposed between the CPU 101 and the main memory 107. The instruction cache 103 operates in accordance with conventional techniques, and is therefore not described here in further detail.

In order to allow the CPU 101 to more quickly retrieve and store data, a data cache 105 is also interposed between the CPU 101 and the main memory 107. The flow of data into and out of the data cache 105 is directed by a data cache controller 109. Much of the operation of the data cache 105 is in accordance with known techniques. Like a conventional cache memory, the data cache 105 is a memory whose cycle time is close to the execution speed of the CPU. Also like a conventional cache memory, the data cache 105 only holds a subset of the information that is contained in the main memory 107. Consequently, the issue of cache misses, as discussed in the BACKGROUND section, is applicable to the data cache 105.

In order to mitigate the effect of cache misses, the architecture of the CPU 101 is capable of performing out-of-order execution of instructions, as described in the BACKGROUND section. Techniques for designing a CPU having this capability are known in the art, and need not be described here in great detail. In order to facilitate an understanding of the invention, a short description of an exemplary embodiment of the CPU 101 follows. It should be understood that the inventive principles described here do not require that the CPU 101 take on the particular form described below, but can instead be applied to any processor architecture that employs out-of-order execution of instructions.

The exemplary CPU 101 includes instruction fetch logic 111 for determining the address of a next instruction and for fetching that instruction. The instruction set of

-8-

the exemplary CPU 101 is designed to utilize operands that are stored in any of a number of addressable working registers 113. Therefore, before an instruction utilizing a source operand can be carried out, that operand must first
5 be loaded into the working register from the data cache 105 by way of operand fetch logic 115, an arithmetic logic unit (ALU) 117, and a re-order buffer 119. Of course, the invention is by no means limited to use only in reduced instruction set computers (RISC) such as the one described
10 here. To the contrary, it may easily be applied to complex instruction set computers (CISC) in which, among other differences, instructions are provided which allow the ALU to operate on data received directly from an addressed memory location without having to first load that data into
15 a working register.

Together, the operand fetch logic 115 and the re-order buffer 119 provide out-of-order execution capability. The operand fetch logic 115 is responsible for assembling source operands which are to be applied to respective inputs
20 of the ALU 117. To facilitate this task, the operand fetch logic 115 includes a number of reservation stations for storing either the source data itself or a sequence number (assigned to the as-yet-unexecuted instruction) which indicates that the source data has not yet been loaded into
25 the working register 113.

The re-order buffer 119 is responsible for assembling destination data, and for keeping track of whether or not a given instruction has been executed, as well as dependencies between instructions. Thus, when an
30 instruction is received, it is assigned a tag number and put into the re-order buffer 119. All previously stored instructions in the re-order buffer 119 are then examined to determine whether any of the sources of the new instruction match the destinations of any other of the previous
35 instructions. If there is a match, then there is a dependency which prevents the new instruction from being

-9-

immediately executed (i.e., the new instruction requires a source operand that has not yet been generated by an earlier instruction). As a result, the sequence number of the new instruction would be placed in the operand fetch logic 115 instead of the source data itself.

In this embodiment, the re-order buffer 119 has space for storing the destination data itself. Consequently, when the destination data becomes available, it is loaded into the re-order buffer 119, and an associated status flag is set to indicate that the data is valid. In order to speed up instruction execution when the destination data is also source data of a subsequent instruction, the destination data may be routed directly from the re-order buffer 119 to the appropriate reservation station in the operand fetch logic 115, rather than requiring that the data first be written to the working register 113 and then moving the data from the working register 113 to the operand fetch logic 115. To speed instruction execution even further, the exemplary CPU 101 further includes a path from the output of the ALU 117 to the input of the operand fetch logic 115, in order to permit newly generated source data to be loaded even sooner.

The fact that generated data has been stored into the re-order buffer 119 is transparent to the executing program, which is only concerned with storing results into one of the working registers 113 or into the data cache 105 (or main memory 107). Consequently, the data must eventually be moved from the re-order buffer 119 into the destination indicated by the corresponding instruction. The act of moving generated data from the re-order buffer 119 to the appropriate destination is called "committing", "retiring" or "writing-back" the data. The architecture of the exemplary CPU 101 adopts the strategy of committing generated data only after all of the data associated with earlier instructions have been committed.

-10-

Despite the inclusion of out-of-order execution capability, the occurrence of a data cache miss still imposes an overhead penalty which slows the effective execution speed of the CPU. 101. Therefore, in order to
5 further increase the effective execution speed of the CPU 101, the data cache controller 109 additionally employs the below-described inventive data replacement techniques whenever a cache miss occurs. Rather than focussing on achieving a high hit ratio, the inventive cache replacement
10 policy is based on reducing the execution speed penalty that will result when data misses occur.

In one aspect of the invention, this objective is achieved by recognizing that data may be classified as one of three categories:

15 Type-1 data may be pure information that is simply written back to an address without being used as an input variable for calculation of any other data item. An example of a type-1 data item is data that is retrieved from memory as part of an editor's copy and paste operation. Another
20 example of a type-1 data item is a parameter associated with a procedure written in C without using pointers to the variables. In such a case, invocation of the procedure will cause the variable to be copied first to the stack, and then to a local variable with the same contents.

25 Type-2 data may be information data. This data is used as an input variable for calculation in a certain number of instructions before being written back to an address.

Type-3 data may be control data. Control data
30 either directly or indirectly determines the operation of a conditional branch instruction. (Control data operates indirectly when it is used as an input variable in any number of calculation instructions for determining a resultant data item which is then used as the predicate for
35 determining the operation of a conditional branch instruction.)

-11-

In the event of a cache miss, the type-1 data results in almost no execution speed penalty so long as the CPU 101 is provided with hardware that allows out-of-order execution. For this purpose, the queues for holding pending
5 instructions and data should be long enough to keep the execution units in the CPU 101 occupied while the missing data is being retrieved.

There is a greater execution speed penalty resulting from a cache miss associated with type-2 data, and
10 the extent of this penalty is a function of how many instructions depend, either directly or indirectly, on the missing data. (Indirect dependence occurs when an instruction requires, as an input, not the missing data itself, but rather the result of a computation that is a
15 function of the missing data.)

In theory, it is possible to utilize an out-of-order execution strategy when a cache miss involves type-2 data. However, in practice, designing a queue that is long enough to hold pending instructions and data is problematic
20 because the number of dependent instructions that need to remain pending might be very large.

The biggest performance problem occurs with respect to cache misses involving type-3 data. If a conditional branch cannot be evaluated, then the CPU 101
25 does not know which instructions to fetch. In the worst case, the instruction that is actually needed is not available in the instruction cache 103. As a result, the CPU 101 will stall during the instruction fetch, thereby severely hurting the effective execution speed.

30 In accordance with one aspect of the invention, the cache replacement policy is designed first to maximize the presence of type-3 data items in the data cache 105. In another aspect of the invention, the cache replacement policy gives type-2 data items priority over type-1 data
35 items in the data cache 105.

-12-

A data item's cache replacement priority is provided by associating that data item with a priority value. For example, priority values may be defined such that the higher the value, the higher the cache miss penalty. Of course, one may easily come up with alternative definitions that are nonetheless in accordance with the inventive principles taught here, such as by letting lower priority values indicate a higher cache miss penalty.

An exemplary embodiment of these aspects of the invention will now be described with reference to the flow chart depicted in FIG. 2. Whenever a data cache miss occurs, the missing data item is retrieved from the main memory 107 and stored into the data cache 105 (step 201). Next, an initial priority value is determined as a function of the penalty that would result if this data item were not present in the data cache 105 at the time that an attempt were made to access it (step 203). The priority value is preferably a function of the present status of queues and other hardware resources in the computer. The adjusted value is then associated with the data item (step 205).

The determination of an initial priority value (step 203) will now be described in greater detail. The priority value of a data item is preferably determined when an attempted access of that data item results in a cache miss. When a cache miss occurs, out-of-order execution of subsequent instructions begins. When instructions are encountered that require the presence of the missing data item, those instructions are stored in a queue until they can be executed. For example, consider the following program segment:

```
load r1, #1;          (r1:="address 1")
add r2, r1, r3;        (r2:=r1+r3)
add r4, r2, r5;        (r4:=r2+r5)
```

In the above example, not only is the initial load instruction dependent on the presence of the data item at address 1 in the data cache 105, but also the subsequent two

-13-

add instructions. This is because the first add instruction requires that r1 have the appropriate value in it, which happens only after the initial load instruction is executed. Similarly, the second add instruction cannot be executed
5 until the contents of r2 are determined. Consequently, in the event of a cache miss associated with "address 1" in the data cache 105, all three of the above instructions will be stored in a queue until the missing data becomes available.

The queue in which instructions are stored is used
10 to determine the priority value of the missing data. If there are dependent conditional branches, then the priority value is very high. If there are data dependencies as in the above example, then the priority value can be set to a medium level. If there are not dependencies, then the
15 priority value of the missing data is low. The priority values might, therefore, be set to 3, 2 and 1, respectively.

A simple approach to determining a priority value is to assume that all executing instructions are directly or indirectly dependent on the missing data. The value in this
20 case is $V = N1 \times V1 + N2 \times V2$, where N1 is the number of non-branch instructions that have not yet been executed, V1 is the penalty value assigned to non-branch instructions, N2 is the number of conditional branches that have not yet been executed, and V2 is the penalty value assigned to
25 conditional branch instructions.

It will be observed that the simple approach described above may result in overly high priority values if some or all of the non-executed instructions are dependent only on a data item other than the one for which the penalty
30 value is being calculated. Also, if there are two data fetches in a code segment, each one resulting in a cache miss, the simple priority value calculation approach will result in high priority values being assigned to each, even if the first data item should accurately have a low priority
35 value and the second data item should have a high priority value.

-14-

In view of the possibility for overestimating priority values under some circumstances, one might perform a more sophisticated analysis of the contents of the instruction buffers in the re-order buffer in order to count only those non-executed instructions that actually have a dependence on the subject data item. However, the overhead associated with this analysis makes it less attractive than the simple approach described above.

Looking now at a more complete example of a priority value calculation, suppose the following code segment is to be executed by the CPU 101:

	<u>Instruction</u>	<u>Comment</u>
	load r0, #5;	(r0 := "address 5")
	jump_on_equal r0, r1;	(jump if r0 = r1)
15	add r6, r3, r0;	(r6 := r3 + r0)
	add r7, r3, r2;	(r7 := r3 + r2)
	add r8, r7, r9;	(r8 := r7 + r9)
	add r9, r10, r11;	(r9 := r10, r11)

where r_x represents working register x.

After these instructions are fetched, the contents of the re-order buffer 119 might look as in Table 1:

-15-

Table 1

	Pos	Seq no	opcode	dest	data	data valid	occupied
5	7						0
	6						0
	5	12	add	r9	1234	1	1
	4	11	add	r8	2344	1	1
	3	10	add	r7	2342	1	1
	2	9	add	r6	-	0	1
	1	8	je	-	-	0	1
	10	0	7	load	r0	-	0

Suppose the "load r0, #5" resulted in a cache miss, thereby requiring that the desired data be fetched from main memory 107. Then, by the time the data from memory address 5 arrives there are one conditional jump (at position 1) and one add (at position 2) which are dependent on the newly arriving data. The "data valid = 1" located at positions 3, 4 and 5 indicates that those add instructions have been executed, and are not dependent on any other instructions. If the penalty for conditional jumps is 5 (i.e., $V2 = 5$) and for other instructions is 1 (i.e., $V1 = 1$), then the total priority value that will be assigned to the address 5 data will be $1 \times 5 + 1 \times 1 = 6$.

Note that, in this case, all of the instructions after the conditional jump instruction were speculatively fetched and executed based on the predicted jump result being jump not taken. If the jump result is mispredicted, then all of the instructions after the conditional jump must be flushed, and the correct instruction path fetched and executed.

-16-

Techniques for determining a priority value for each data cache entry have been described. As indicated in FIG. 2, it is necessary to associate the priority value with the cache entry. This may be done in any of a variety of ways. One of these will now be described with reference to FIG. 3, which depicts an exemplary embodiment of a data cache 105 for use with the above-described inventive techniques is shown. The data cache 105 is capable of storing N data items at addresses designated from 0 to N-1. At an exemplary address, m, a data item 301 is stored. The data cache 105 is constructed such that for each location that stores a data item, there is a corresponding location for storing an associated priority value. Thus, for the exemplary location m, the stored data item 301 is associated with a stored priority value 303. Techniques for constructing a cache memory having this capability are well-known in the art, and are not described here in greater detail. See, for example, John L. Hennes, Computer Architecture, A Quantitative Approach, (1990) which is hereby incorporated herein by reference.

In another aspect of the invention, a strategy for replacing data in a cache memory will now be described with respect to the flow chart of FIG. 4. In this exemplary embodiment, it is assumed that each entry in a data cache 105 has an associated priority value as described above.

When an attempt is made to read an entry from the data cache 105, a determination is made regarding whether the requested data item is presently stored in the data cache 105 (decision block 401). If the item is found ("no" path out of decision block 401), then the item is simply retrieved from the data cache 105 (step 403), and processing continues.

If a cache miss occurs ("yes" path out of decision block 401), then it is necessary to remove an entry from the data cache 105 in order to make room for the requested data item. This is accomplished by locating an entry having a

-17-

priority value that is lowest among all of the priority values 303 stored in the data cache 105 (step 405). In the event that more than one entry has a lowest priority value, any technique may be used for deciding which one to remove from the data cache 105: random selection, use of a pointer, or any other technique. Then, the requested (i.e., "missing") data item is retrieved from the main memory 107 (step 407), and the retrieved data item is stored into the data cache 105 at the address that was identified in step 405. By removing an entry having a lowest priority value, the inventive cache replacement strategy reduces the likelihood that a subsequent cache miss associated with the removed entry will occur (i.e., the priority value was low because the removed data item had not been accessed for a long time), or alternatively makes it less likely that a subsequent cache miss associated with the removed entry will result in any substantial execution speed penalty (i.e., because out-of-order execution of subsequent instructions can proceed while waiting for retrieval of the requested data item from main memory 107).

The above-described data cache replacement strategy, including the determination of priority values for each cache entry, increases the likelihood that the data in the cache will be the kind of data that would cause a substantial execution speed penalty if it had to be retrieved from the main memory 107. That is, if relative priority values are 1, 2 and 3 for Type-1, -2, and -3 data, respectively, it is more likely that Type-3 data will be in the cache because the Type-1 and Type-2 are more likely to be replaced. However, the above-described strategy creates the possibility that a data item will stay in the data cache 105 indefinitely, even if that data item is never used again by the CPU 101. This is because so long as there is a lower-priority cache entry, such a data item will never be replaced in the data cache 105. The presence of such a cache entry (i.e., one that is never replaced even though it

-18-

is never again required as a source operand) is undesirable because it reduces the cache hit/miss ratio.

This problem is addressed in yet another aspect of the invention, by allowing the priority value to further be a function of how much time has passed since the data item was last accessed. In this way, the principle of temporal locality may also be employed to increase the hit ratio of the cache, while still giving preference to data types having higher execution speed penalties. In an exemplary embodiment, the priority value 303 of each cache entry is decremented by a predetermined amount every T seconds. As a result, even those cache entries having initially-high priority values will eventually have low-enough priority values to make them contenders for replacement when a new cache entry is to be stored.

In yet another aspect of the invention, the strategy in which priority values are periodically decremented is supplemented by re-initializing or alternatively incrementing the priority value 303 whenever the cache entry is accessed. In one embodiment, the priority value 303 may be increased by the associated penalty value whenever the cache entry is accessed. In this way, the priority value takes into account not only the cost of one cache miss, but also the frequency with which this data entry is accessed. For example, assume that a first data item costs only one clock cycle of CPU stall time in the event of a cache miss, and that this data item is accessed ten times as frequently as a second data item costing ten clock cycles of CPU stall time in the event of a cache miss. In this case, the value of the first data item is equal to the value of the second data item. Thus, by reinitializing or incrementing the priority value every time a data item is accessed, the replacement strategy takes into account the effect of frequent access.

If the strategy is employed in which the priority value 303 is reinitialized with every access of the cache

-19-

entry, it is necessary to keep track not only of the current priority value 303 for each cache entry, but also the initial priority value. This may be accomplished by organizing a data cache 105' as shown in FIG. 5. The data
5 cache 105' is substantially identical to the data cache 105 shown in FIG. 3 except for the addition of an initial priority value field 501 associated with each cache entry.

The discussion will now focus on the determination of relative penalty values (i.e., V1 and V2) for different
10 kinds of instructions. Suppose a cache miss resulting from an attempt to retrieve a first data item results in ten clock cycles of CPU stall time, and that a cache miss resulting from an attempt to retrieve a second data item causes only one clock cycle of CPU stall time because of the
15 possibility of out-of-order execution. It follows that, given a choice between the two, the first data item should be retained in the data cache 105 rather than the second data item in order to avoid an expensive cache miss. In this sense, the first data item is more "valuable" than the
20 second data item.

Consider a third data item that, because of very effective out-of-order execution, does not cause any CPU stall time when a cache miss occurs. Such a data item need not be stored in the data cache at all; that is, the third
25 data item has little, if any, value.

In view of the above, it can be stated that the more valuable a data item is, the longer it should be retained in the data cache 105. A good rule of thumb is to retain a data item in proportion to its value. For example,
30 the first data item costing ten clock cycles of CPU stall time in the event of a cache miss should be retained in the data cache 105 ten times longer than the second data item costing only one clock cycle of CPU stall time (assuming that neither data item is accessed during this time period).
35 This may easily be accomplished by setting the penalty value of the first data item equal to 10 (i.e., V1=10), setting

-20-

the penalty value of the second data item equal to 1 (i.e., $V_2=1$), determining a priority value 303 for each, as described above, and then adjusting (e.g., decreasing) the magnitude of each priority value at a predetermined frequency, while also reloading the value every time the data item is accessed.

The above discussion of determining penalty values (e.g., V_1 , V_2) is premised on the availability of measurements or estimates of CPU stall times for different types of instructions. In the illustrative embodiment, the penalty values associated with different data types is fixed for each determination of an initial priority value 501. When type-1 or type-2 data items are considered, it is a relatively simple matter to come up with a value that represents an average CPU stall time, because the execution speed of these types of instructions (e.g., load, add, logical AND) is easily determined.

With respect to type-3 data items, however, the task becomes more difficult because the execution speed is a function of whether or not the branch is taken, and whether or not prefetched instructions (based on a prediction of whether or not the branch would be taken) need to be flushed from the instruction cache 103 and new instructions fetched and executed. Those having ordinary skill in the art will recognize that an average penalty value for type-3 data can be determined based on knowledge about the processor speed, the access time to the instruction memory, the cache hit ratio and the relative accuracy of the branch prediction algorithm (e.g., 90% accurate prediction).

In an alternative embodiment, the penalty values associated with different data types may be dynamically determined for each individual data item. With respect to type-1 or type-2 data items, the dynamically determined values will most likely be constant (or nearly constant), since the respective execution speeds of these types of instructions (e.g., load, add, logical AND) do not generally

-21-

change from time to time. (In some architectures, the execution speed of instructions such as multiply and divide may be somewhat data dependent.)

In this alternative embodiment, the penalty values associated with type-3 data can be dynamically determined by allowing any pending conditional branch instructions to execute after the missing data arrives, and then actually measuring to see what the delay time turns out to be (e.g., seeing whether prefetched instructions need to be flushed because of an inaccurate branch prediction, and if so, how many instructions were flushed). Of course, the overhead associated with this dynamic determination of each penalty value is considerably more than that associated with the use of a fixed estimate of penalty value. Also, determining the penalty value dynamically delays a final determination of initial priority value 501. These factors should be taken into consideration in deciding which approach (i.e., static versus dynamic penalty value determination) should be used for any particular system.

A strategy for maintaining priority values in accordance with the above-described principles will now be described with reference to the flow chart of FIG. 6. It is assumed that a data item (such as the data item 301) and its corresponding current priority value 303 and initial priority value 501 are stored in a data cache 105'. In this example, it is further assumed that higher priority values indicate a higher priority for retaining the data item in the data cache 105'.

The exemplary priority value maintenance strategy employs the principle of temporal locality in order to maintain a high hit ratio in the data cache 105'. Accordingly, if it is determined that it has been more than T seconds since a data item was last accessed ("yes" path out of decision block 601), the priority value that is associated with the data item is decreased by some predetermined amount (step 603). In a preferred embodiment,

-22-

the predetermined amount is a lowest possible amount, such an amount equal to one.

If it is determined that the data item has been recently accessed ("yes" path out of decision block 605), then its associated priority value is increased or alternatively set to the initial calculated priority value (step 607). The reason for doing this is because the recent access indicates a higher likelihood that this data item will be accessed again in the near future. Consequently, the increase (or reinitialization) of the priority value will help ensure that the data item will not be removed, but instead will remain in the data cache 105' for the near term. In a preferred embodiment, step 607 includes distinguishing between read accesses and write accesses, so that every read access causes the priority value to be increased by the initial calculated priority value (e.g., an amount ranging between one and approximately eight), whereas every write access causes the priority value to be increased by a lowest possible value (e.g., an amount equal to one).

The description so far has been premised on the data cache 105 being controlled in a manner such that each cache miss causes only one data item to be retrieved from the main memory 107 and stored into the data cache 105. However, it is more typical that data caches are designed such that each cache miss causes a number (e.g., 8, 16) of data items (referred to herein as a "cache entry") to be retrieved from the main memory 107 and stored into the data cache 105. The cache entry usually comprises not only the "missing" data item (i.e., the contents of the main memory address location whose absence in the data cache 105 caused the cache miss), but also those data items whose addresses are near that of the "missing" data item. The rationale for this strategy is to improve the data cache's hit ratio in reliance on a principle known as "spatial locality" (or "locality of reference"). This principle, which is based on observations of how programs are typically designed, asserts

-23-

that if a memory access presently reads data from, say, location x, then there is an increased probability that a subsequent memory access will attempt to read data from a memory address location that is near location x. Thus, by
5 retrieving a cache entry comprising not only the data item from location x but also those data items stored at locations close to x, it is more likely that a subsequent attempt to access the data cache 105 will result in a cache hit.

10 In another embodiment, the inventive principles described above are applied to a data cache whose data cache controller 109 retrieves cache entries (as opposed to single data items) whenever a cache miss occurs. An exemplary data cache 105'' adapted for this purpose is shown in FIG. 7.
15 The current priority value 303 and initial priority value 501 are determined and utilized in the same manner as described above. The difference between the data cache 105'' and previously described embodiments lies in the fact that each pair of current priority value 303 and initial
20 priority value 501 is associated with an entire cache entry 701, rather than being uniquely associated with just one data item 301. Determination of the initial priority value 501 is based on whichever one of the data items 301 in the cache entry 701 caused the cache miss to occur, thereby
25 causing the cache entry 701 to be retrieved from the main memory 107. An assumption is made that other data items 301 in the cache entry 701 are of a similar data type, and so would have similar priority values. In another aspect of this embodiment, whenever any of the data items 301 in the
30 cache entry 701 is accessed, the value of the associated current priority value 303 may be adjusted (e.g., reinitialized or increased) as described above with respect to FIG. 6. In yet another aspect of the invention, the current priority value 303 may be adjusted to a value
35 indicating lower priority if none of the data items 301 in the cache entry 701 have been accessed during any particular

-24-

period of time. In all other respects, the principles involved with this embodiment are the same as those described above, and therefore need not be described again.

The invention has been described with reference to
5 a particular embodiment. However, it will be readily apparent to those skilled in the art that it is possible to embody the invention in specific forms other than those of the preferred embodiment described above. This may be done without departing from the spirit of the invention. The
10 preferred embodiment is merely illustrative and should not be considered restrictive in any way. The scope of the invention is given by the appended claims, rather than the preceding description, and all variations and equivalents which fall within the range of the claims are intended to be
15 embraced therein.

-25-

WHAT IS CLAIMED IS:

1. In a computer system having a central processing unit (CPU), a cache memory and a main memory, a method of
5 storing a data item into the cache memory comprising the steps of:
initiating retrieval of the data item from the main memory in response to at least partial execution of a first instruction;
10 determining a priority value to be associated with the data item, wherein the priority value is an estimate of how much CPU stall time will occur if an attempt is made to retrieve the data item from the cache memory when the data item is not stored in the cache memory; and
15 completing retrieval of the data item and storing both the data item and the priority value in the cache memory.
2. The method of claim 1, wherein the step of
20 determining the priority value comprises determining how many other instructions have been fetched and stored in a buffer memory between the time of initiating retrieval of the data item and completing retrieval of the data item, wherein execution of the other instructions is dependent on
25 completing retrieval of the data item.
3. The method of claim 2, wherein the step of
determining the priority value further comprises adjusting the priority value by an amount based on whether any of the
30 other instructions is a conditional branch instruction.
4. The method of claim 2, wherein the step of
determining the priority value further comprises adjusting the priority value by an amount based on whether any of the
35 other instructions is a non-branch instruction.

-26-

5. The method of claim 1, wherein the step of determining the priority value occurs after retrieval of the data item has been completed.

5 6. A cache data replacement method for use in a computer system having a central processing unit (CPU), a cache memory and a main memory, wherein the cache memory has a plurality of data items stored therein, the cache data replacement method comprising the steps of:

10 associating a priority value with each of the stored data items, wherein for each data item, the priority value is an estimate of how much CPU stall time will occur if an attempt is made to retrieve the data item from the cache memory when the data item is not stored in the cache
15 memory;

analyzing the priority values to determine a lowest priority value;

selecting one of the data items that has the lowest priority value; and

20 replacing the selected data item with a replacement data item.

7. The cache data replacement method of claim 6, wherein the step of associating a priority value with each
25 of the stored data items comprises the step of determining, for each of the data items, how many other instructions have been fetched and stored in a buffer memory between a time interval defined by initiation and completion of retrieval of the data item from the main memory, wherein execution of
30 the other instructions is dependent on completing retrieval of the data item.

8. The cache data replacement method of claim 7, wherein the step of determining, for each of the data items,
35 how many other instructions have been fetched and stored in the buffer memory between the time interval further

-27-

comprises adjusting the priority value by an amount based on whether any of the other instructions is a conditional branch instruction.

5 9. The cache data replacement method of claim 7,
 wherein the step of determining, for each of the data items,
 how many other instructions have been fetched and stored in
 the buffer memory between the time interval, further
 comprises adjusting the priority value by an amount based on
10 whether any of the other instructions is a non-branch
 instruction.

 10. The cache data replacement method of claim 6,
 further comprising the step of:
15 periodically adjusting each of the priority values
 by a predetermined amount that results in a lower priority
 value.

 11. The cache data replacement method of claim 6,
20 further comprising the step of:
 for each of the priority values, adjusting the
 priority value by a first predetermined amount in response
 to a read access of the associated data item, wherein the
 adjustment results in a higher priority value.

25 12. The cache data replacement method of claim 11,
 further comprising the step of:
 for each of the priority values, adjusting the
 priority value by a second predetermined amount in response
30 to a write access of the associated data item, wherein the
 adjustment results in a higher priority value.

 13. The cache data replacement method of claim 12,
 wherein adjustment by the first predetermined amount results
35 in a higher priority value than adjustment by the second
 predetermined amount.

-28-

14. The cache data replacement method of claim 6, further comprising the step of:

for each of the priority values, setting the priority value to an initial priority value in response to a read access of the associated data item.

15. In a computer system having a central processing unit (CPU), a cache memory and a main memory, an apparatus for storing a data item into the cache memory comprising:

means for initiating retrieval of the data item from the main memory in response to at least partial execution of a first instruction;

means for determining a priority value to be associated with the data item, wherein the priority value is an estimate of how much CPU stall time will occur if an attempt is made to retrieve the data item from the cache memory when the data item is not stored in the cache memory; and

means for completing retrieval of the data item and storing both the data item and the priority value in the cache memory.

16. The apparatus of claim 15, wherein the means for determining the priority value comprises means for determining how many other instructions have been fetched and stored in a buffer memory between the time of initiating retrieval of the data item and completing retrieval of the data item, wherein execution of the other instructions is dependent on completing retrieval of the data item.

17. The apparatus of claim 16, wherein the means for determining the priority value further comprises means for adjusting the priority value by an amount based on whether any of the other instructions is a conditional branch instruction.

-29-

18. The apparatus of claim 16, wherein the means for determining the priority value further comprises means for adjusting the priority value by an amount based on whether any of the other instructions is a non-branch instruction.

5

19. The apparatus of claim 15, wherein the means for determining the priority value makes the determination after retrieval of the data item has been completed.

10 20. A cache data replacement apparatus for use in a computer system having a central processing unit (CPU), a cache memory and a main memory, wherein the cache memory has a plurality of data items stored therein, the cache data replacement apparatus comprising:

15 means for associating a priority value with each of the stored data items, wherein for each data item, the priority value is an estimate of how much CPU stall time will occur if an attempt is made to retrieve the data item from the cache memory when the data item is not stored in
20 the cache memory;

means for analyzing the priority values to determine a lowest priority value;

means for selecting one of the data items that has the lowest priority value; and

25 means for replacing the selected data item with a replacement data item.

21. The cache data replacement apparatus of claim 20, wherein the means for associating a priority value with each
30 of the stored data items comprises means for determining, for each of the data items, how many other instructions have been fetched and stored in a buffer memory between a time interval defined by initiation and completion of retrieval of the data item from the main memory, wherein execution of
35 the other instructions is dependent on completing retrieval of the data item.

-30-

22. The cache data replacement apparatus of claim 21, wherein the means for determining, for each of the data items, how many other instructions have been fetched and stored in the buffer memory between the time interval
5 further comprises means for adjusting the priority value by an amount based on whether any of the other instructions is a conditional branch instruction.

23. The cache data replacement apparatus of claim 21,
10 wherein the means for determining, for each of the data items, how many other instructions have been fetched and stored in the buffer memory between the time interval, further comprises means for adjusting the priority value by an amount based on whether any of the other instructions is
15 a non-branch instruction.

24. The cache data replacement apparatus of claim 20, further comprising:
means for periodically adjusting each of the
20 priority values by a predetermined amount that results in a lower priority value.

25. The cache data replacement apparatus of claim 20, further comprising:
25 means for adjusting the priority value of an associated data item by a first predetermined amount in response to a read access of the associated data item, wherein the adjustment results in a higher priority value.

30 26. The cache data replacement apparatus of claim 25, further comprising:
means for adjusting the priority value of an associated data item by a second predetermined amount in response to a write access of the associated data item,
35 wherein the adjustment results in a higher priority value.

-31-

27. The cache data replacement apparatus of claim 26, wherein adjustment by the first predetermined amount results in a higher priority value than adjustment by the second predetermined amount.

5

28. The cache data replacement apparatus of claim 20, further comprising:

means for setting the priority value of an associated data item to an initial priority value in response to a read access of the associated data item.

10




FIG. 1

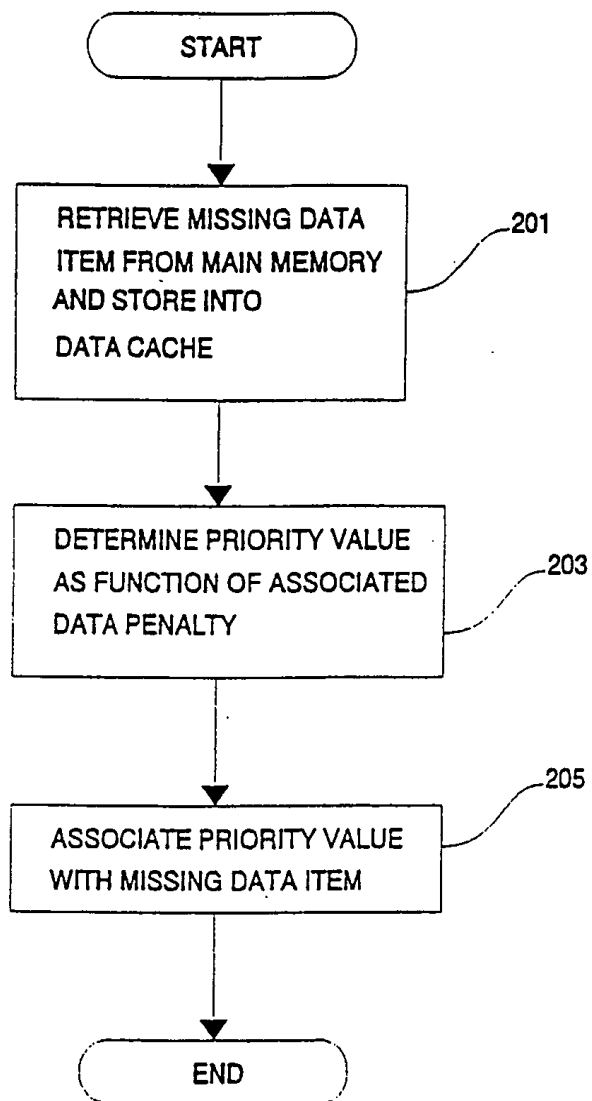


FIG. 2




FIG. 3

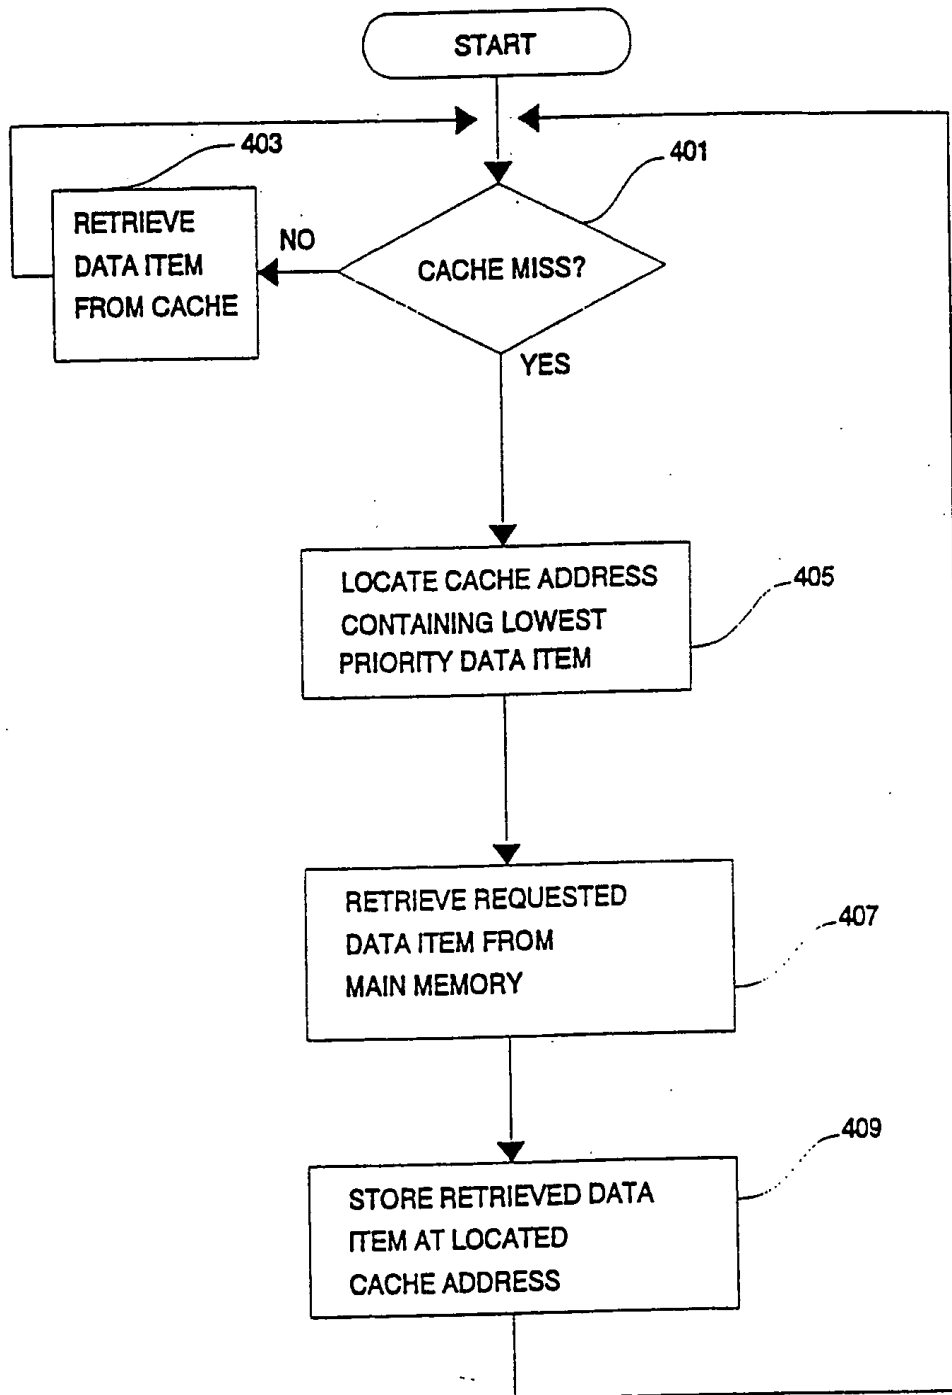


FIG. 4




FIG. 5

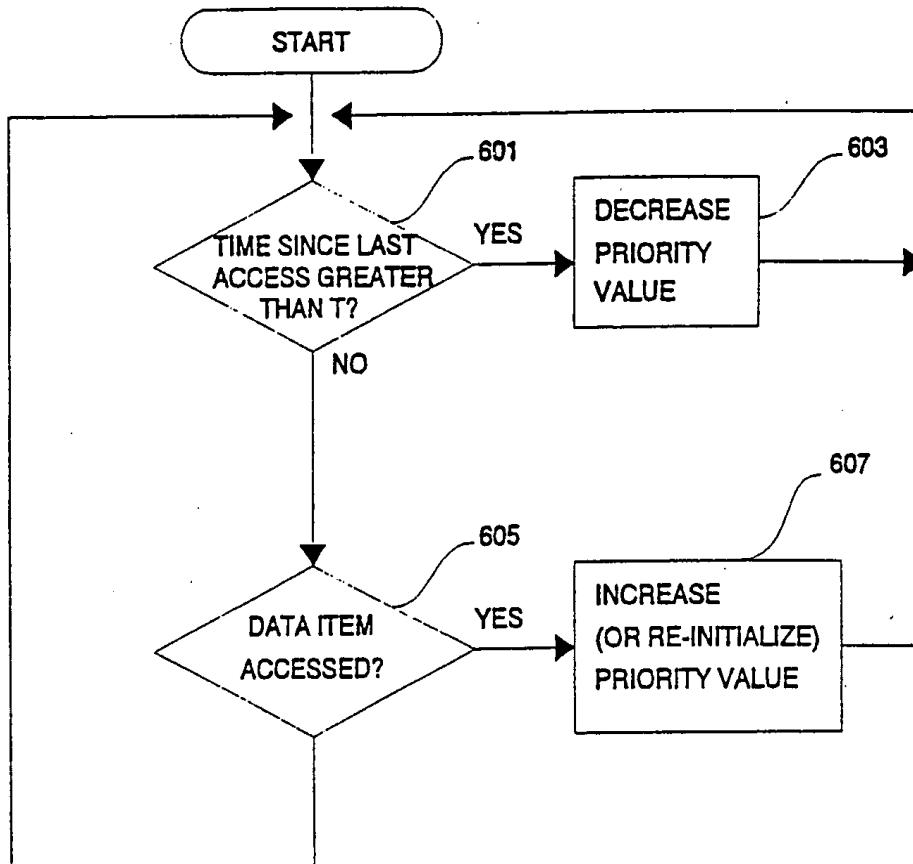


FIG. 6




FIG. 7

INTERNATIONAL SEARCH REPORT

International Application No

PCT/SE 98/00454

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F12/12 G06F12/08

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	"ALGORITHM FOR INSTRUCTION CACHE ARBITRATION AMONG MULTIPLE INSTRUCTION STREAMS" IBM TECHNICAL DISCLOSURE BULLETIN, vol. 39, no. 8, August 1996, pages 113-116, XP000638156 see page 115, paragraph 2 - page 116, paragraph 4; figures 1,2 ---	1,15
A	CHIA ET AL.: "Replacement algorithm using priority class structure" IBM TECHNICAL DISCLOSURE BULLETIN., vol. 15, no. 12, May 1973, NEW YORK US, pages 3803-3805, XP002067195 see the whole document --- -/--	6,20

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "Z" document member of the same patent family

Date of the actual completion of the international search

5 June 1998

Date of mailing of the international search report

22/06/1998

Name and mailing address of the ISA
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Ledrut, P

INTERNATIONAL SEARCH REPORT

Intern. Appl. No.

PCT/SE 98/00454

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>BAE J -H ET AL: "A SUPPLEMENTARY SCHEME FOR REDUCING CACHE ACCESS TIME" IEICE TRANSACTIONS ON INFORMATION AND SYSTEMS, vol. E79-D, no. 4, 1 April 1996, pages 385-387, XP000592767 see page 385, right-hand column, paragraph 4 - page 386, left-hand column, paragraph 2; figure 2</p> <p>-----</p>	6,20